arXiv:1607.08176v1 [cs.DS] 27 Jul 2016

Suffix arrays with a twist

Tomasz Kowalski[†], Szymon Grabowski[†],
Kimmo Fredriksson[‡] and Marcin Raniszewski[†]

[†] Lodz University of Technology, Institute of Applied Computer Science,
Al. Politechniki 11, 90–924 Łódź, Poland
{tkowals|sgrabow|mranisz}@kis.p.lodz.pl
[‡] School of Computing, University of Eastern Finland,
P.O.B. 1627, FI-70211 Kuopio, Finland, kimmo.fredriksson@uef.fi

Abstract. The suffix array is a classic full-text index, combining effectiveness with simplicity. We discuss three approaches aiming to improve its efficiency even more: changes to the navigation, data layout and adding extra data. In short, we show that (*i*) how we search for the right interval boundary impacts significantly the overall search speed, (*ii*) a B-tree data layout easily wins over the standard one, (*iii*) the well-known idea of a lookup table for the prefixes of the suffixes can be refined with using compression, (*iv*) caching prefixes of the suffixes in a helper array can pose a(nother) practical space-time tradeoff.

1 Introduction

Everybody knows the suffix array (SA) [4], a simple full-text index data structure capable of finding the *occ* occurrences of a pattern P of length m in $O(m \log n + occ)$ time, where n is the length of the indexed text. The search mechanism consists in two binary searches, for the left and the right boundary of the interval of text suffixes starting with P, in the array of suffix offsets arranged in the lexicographical order of their text content. The performance of the suffix array can serve as a measuring stick for more advanced (e.g., compressed) text indexes [1] and at least for this reason it is important to know how to implement it efficiently and what space-time tradeoffs are possible.

The body of research on engineering the suffix array is surprisingly scarce. Although the basic SA idea can be easily grasped even by high-school students, many design choices from the implementor's point of view are not obvious. Let us pose a few questions: (*i*) Can the binary search strategy be replaced with a faster one, e.g., based on interpolation search? (*ii*) As *occ* is usually small, what is practically the best way to find the right interval boundary once the left boundary is known? (*iii*) Can we change the data layout of suffixes in order to obtain more local memory accesses? (*iv*) How can we augment the suffix array with a moderate amount of extra data, to initially reduce the search interval and/or speed up string comparisons? The answer to some of them is known, yet in this work we are dealing with the mentioned issues in a more systematic way.

2 Ideas and incarnations

The considered ideas are divided into three groups and each of them is described in a separate paragraph. First we discuss non-standard SA traversal strategies. Later we advocate for alternative data layouts, beneficial for the search speed. In the last paragraph some ways to augment the suffix array with extra data, to make the pattern search even faster, are proposed.

Navigating over the suffix array. A textbook alternative to binary search is interpolation search, which performs a number of "guesses" concerning the query's location based on the query value and the assumed distribution of keys. It is well-known that interpolation search over constant-size keys achieves $O(\log \log n)$ expected time not only for the simplest case, i.e., uniformly random distribution [6], yet we are not aware of any published experiments regarding text suffixes. Unfortunately, a straightforward interpretation of string prefixes (which have a lot of duplicates) as integers and standard linear interpolation yielded rather disappointing preliminary results.

Another question concerning the navigation over the SA is how the right interval boundary should be found. We examine two methods: a naive one performs the binary search over the range $left \ldots n$ of suffixes, where $left$ is the position of the least suffix greater or equal to the pattern, and the doubling (galloping) algorithm, which peeks the locations $SA[left + 2^i], i = 0, 1, 2, \ldots$, until it reaches too far and the search continues in the binary manner over the last considered interval. Note that the time complexity of the right interval boundary search improves in this way from $O(m \log n)$ to $O(m \log occ)$.

Linearized k-ary tree data layout. Binary search over a sorted array is equivalent to walking down a path in a complete binary search tree. Schleger et al. [5] noticed that changing the tree layout from binary to k-ary ($k > 2$), together with linearization of the search tree, may be more cache-friendly and also convenient for SIMD processing. In their experiments (Intel Core i7) it achieved a speedup of as much as 3 up to 4.5 for 32-bit numbers and 2 to 2.5 for 64-bit numbers, compared to a plain binary search. This data organization can also be called an (implicit) B-tree layout [3]. We apply the presented idea to the suffix array, which, to our knowledge, has not been tried before. Note that setting the B-tree layout for a suffix array cannot be comparably successful as for, e.g., integers, as the accesses to the text are still at "random" areas.

Augmenting the suffix array. Manber and Myers in their seminal paper [4] presented a nice trick saving several first steps in the binary search: if we know the SA intervals for all the possible first k symbols of the pattern, we can immediately start the binary search in a corresponding interval. We can set k to $\log_\sigma n$, where σ is the alphabet size, with $O(n \log n)$ extra bits of space and constant expected size of the interval. Unfortunately, real texts are far from random, hence in practice, we can use k up to 3 (assuming that text symbols are bytes), which offers a limited (yet, non-negligible) benefit. This idea will be referred in our

experiments as using a lookup table, and more specifically we will denote the lookup table on pairs (resp. triples) of symbols with LUT2 (resp. LUT3).

In the same spirit, Grabowski and Raniszewski [2] use a hash table to store the intervals for all k-symbol strings *occurring in the text*. This can significantly reduce the initial interval for real texts with relatively little extra space.

In this work we first propose a lookup table with keys being concatenations of Huffman codewords for the starting symbols of the text suffixes (Table 1), truncated to a specified length of b bits. Pattern search translates to finding the first b bits of Huffman encoding of the pattern, which is the LUT key, and then following with binary search over a range of suffixes read from the LUT. A correct implementation of this idea requires a reordering of the suffixes in the SA, to avoid nested LUT ranges (other options, like replacing Huffman with Hu–Tucker coding, are also possible but we have not tried them out). Note also that the Huffman-based LUT entries store twice more data (both boundaries of the interval) than in the standard LUTs.

Table 1. Average binary logarithms of the search interval widths for different LUT variants (first three rows: Huffman encoding with 16–24 bits, next two rows: standard 2-/3-byte LUTs). The averages are over 100K patterns taken randomly from the text.

dataset	space (MiB)	dna200	english200	proteins200	xml200
—	—	27.644	27.644	27.644	27.644
LUT-Huff-15b	0.25	14.453	15.530	12.858	17.045
LUT-Huff-19b	4.00	11.070	13.432	9.103	15.898
LUT-Huff-23b	64.00	7.790	11.630	5.617	14.985
LUT2	0.25	23.738	19.502	19.266	18.910
LUT3	64.00	21.820	16.555	15.108	16.756

We also propose mixing the LUT or hash table interval narrowing with the B-tree layout, and also augmenting the search tree with prefixes of the suffixes in several top levels of the B-tree. Copying these text snippets into a helper array is beneficial due to more local memory accesses.

3 Experimental results

The test machine was equipped with a 4-core Intel i7 4790 3.6 GHz CPU and 32 GB of 1600 MHz DDR3 RAM (9-9-9-24), hosting Windows Server 2012 R2. All codes (`https://bitbucket.org/kowallus/sa-search-dev/`) were written in C++ and compiled with 64-bit gcc 4.9.3 with `-O3`. All presented times are averages over 500K patterns taken randomly from the text. The searches were performed over 200 MB datasets from the well-known Pizza & Chili corpus.

In the first experiment we show how the count times are affected by two things: using lookup tables, including the introduced Huffman-based ones (on 15 or 23 bits) and choosing a proper interval's right boundary search (Fig. 1).

The doubling trick reduces the times usually by 20–30% for the standard and LUT2-boosted suffix array, yet the effect is smaller for short patterns, especially for DNA (where short patterns tend to have thousands of occurrences). This can be explained by the relatively small difference between $\log n$ and $\log occ$ in those cases. The Huffman-based LUTs are more efficient than their traditional counterparts (when about the same amount of memory is sacrificed).

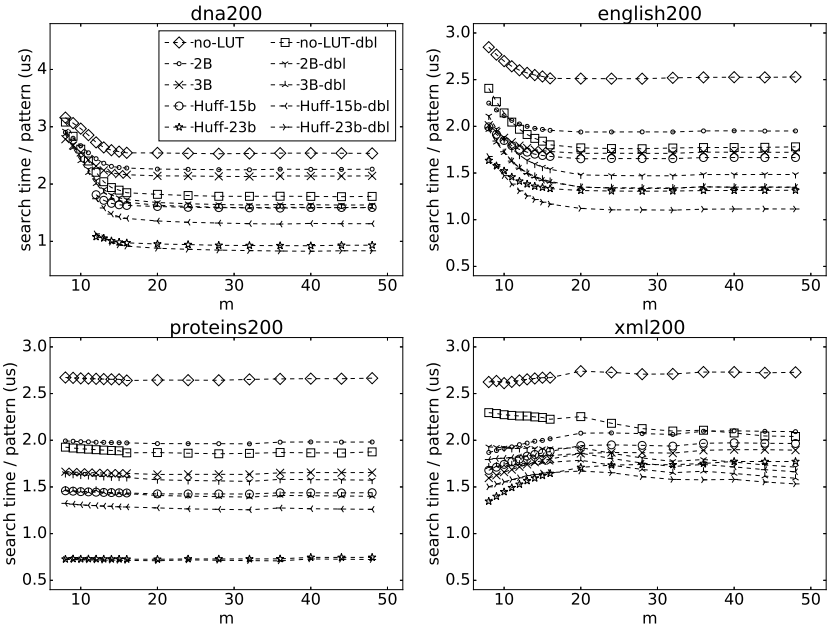

Fig. 1. Count times for the standard suffix array and the SA augmented with a lookup table on pairs or triples of bytes, and on 15 or 23 bits of Huffman codewords, using the standard and the doubling search for finding the right interval boundary.

Fig. 2 shows the impact of the node size B in the B-tree layout on the count times, with varying pattern length. Even $B = 1$ results in a much faster search than with a standard SA (by a factor of 1.7–2.0; cf. also Fig. 1) and growing B helps more, up to $B = 32$ (on all the datasets, $B = 64$ is slightly slower). Still, the speed gap between $B = 1$ and $B = 32$ rarely exceeds 10%.

In Fig. 3 we show how augmenting the SA with various structures reducing the initial search interval affects the query times and the used space. The hash table (HT), based on the xxhash function (`https://code.google.com/p/xxhash/`), stores 8-grams from the text and was tried with two load factors (LF). LUT2 gives a significant boost in a tiny space, yet it is the hash table (LF=0.9) that excels here, speeding up the baseline variant by a factor of 1.5–2.

As our SAs with the B-tree layout can be augmented with prefixes of the suffixes visited in the first steps of the traversal of the tree (i.e., in the top

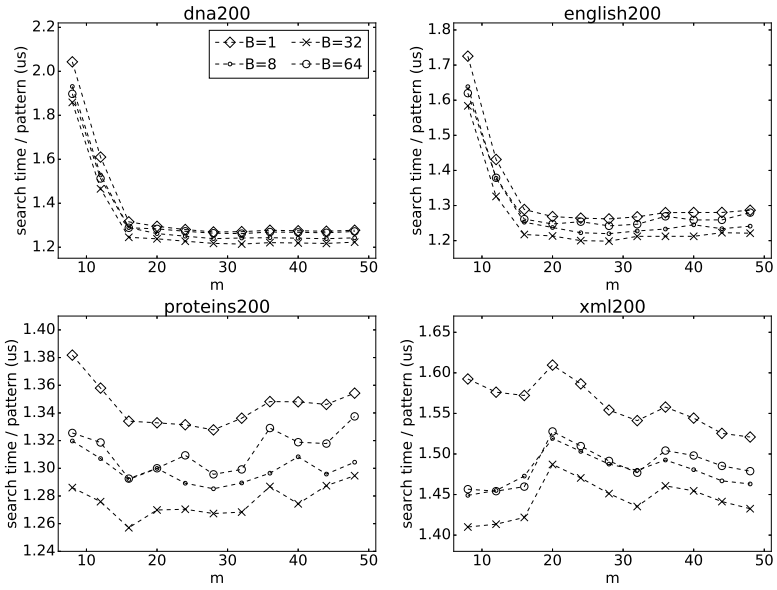

Fig. 2. Count times for the SA with the B-tree layout, for selected node sizes B.

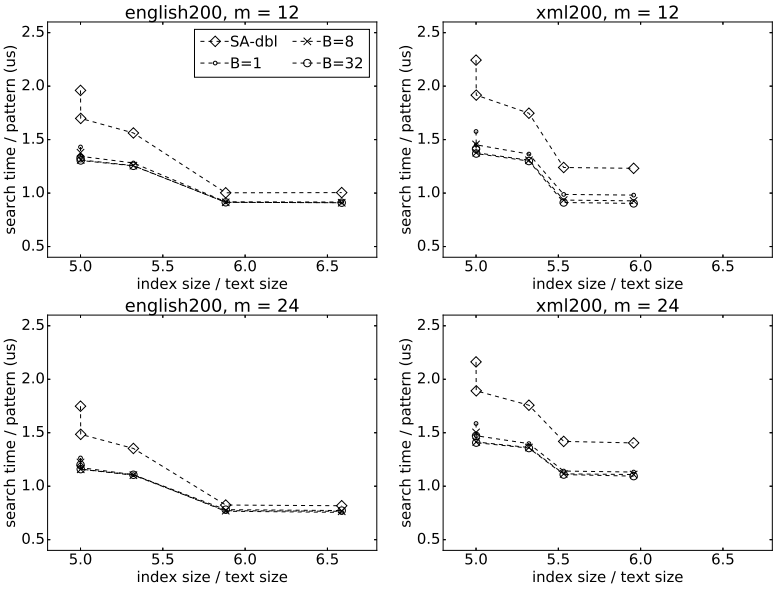

Fig. 3. Augmenting the suffix array. The right interval boundary in the standard SA variant is found with the doubling technique (-dbl). The five points in each series correspond to: no extra data, LUT2, LUT3, HT with LF=0.9, HT with LF=0.5.

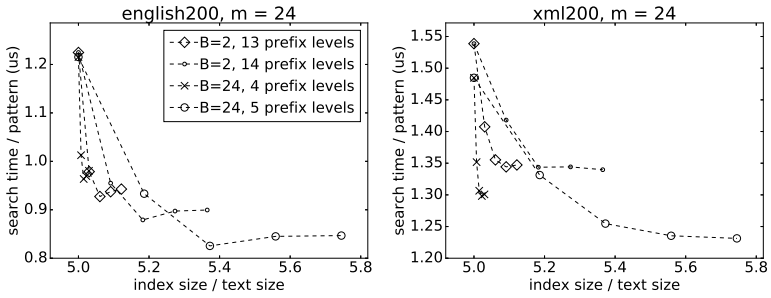
english200, m = 24
xml200, m = 24
B=2, 13 prefix levels
B=2, 14 prefix levels
B=24, 4 prefix levels
B=24, 5 prefix levels
search time / pattern (us)
index size / text size

Fig. 4. Index sizes and count times for the SA with the B-tree layout, when several top levels of the tree store the corresponding suffixes' prefixes of length $\{0, 4, 8, 12, 16\}$.

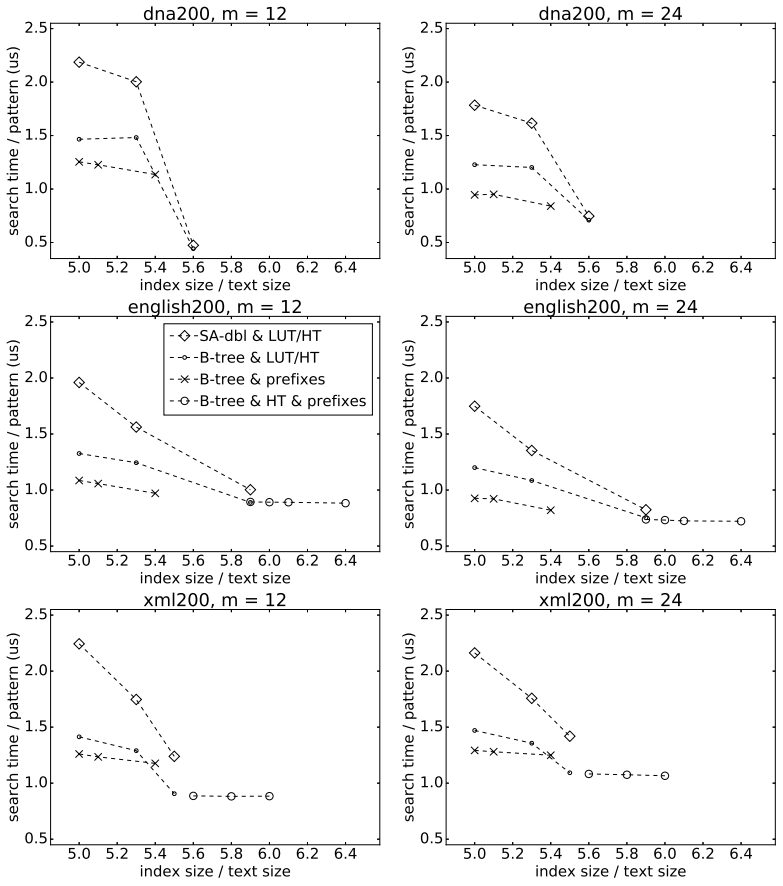
dna200, m = 12
dna200, m = 24
english200, m = 12
english200, m = 24
xml200, m = 12
xml200, m = 24
SA-dbl & LUT/HT
B-tree & LUT/HT
B-tree & prefixes
B-tree & HT & prefixes
search time / pattern (us)
index size / text size

Fig. 5. Index sizes and count times for several SA variants with different layouts and extra data. Successive points in the series are obtained by changing the LUT or hash table component and/or using the prefix copies on varying number of levels in the tree.

levels), we test the impact of the prefix length on the performance and space of the resulting index (Fig. 4). Adding the prefixes gives a noticeable speedup even if they are limited in length to 8 characters and are attached to a few tree levels only, while the space overhead is rather small. Longer prefixes and more levels help less for a much bigger space penalty. Combinations of all ideas presented in the paper are shown in Fig. 5, where the best option is to combine the B-tree layout with LUT/HT (adding prefixes on top of it has a negligible effect). In total, the speed of the standard SA with the standard right interval boundary search was usually improved by a factor exceeding 3 (from 2.6 for xml200 to 3.9 for dna200, for $m = 24$).

Acknowledgment The work was supported by the Polish National Science Centre under the project DEC-2013/09/B/ST6/03117 (the first two and the fourth author).

References

1. P. Ferragina, R. González, G. Navarro, and R. Venturini. Compressed text indexes: From theory to practice. *ACM Journal of Experimental Algorithmics*, 13:article 12, 2009. 30 pages.
2. S. Grabowski and M. Raniszewski. Two simple full-text indexes based on the suffix array. In *Proc. PSC*, pages 179–191, Czech Technical University in Prague, Czech Republic, 2014.
3. P. Khuong and P. Morin. Array layouts for comparison-based searching. *CoRR*, abs/1509.05053, 2015.
4. U. Manber and G. Myers. Suffix arrays: a new method for on-line string searches. In *Proc. SODA*, pages 319–327. SIAM, 1990.
5. B. Schlegel, R. Gemulla, and W. Lehner. K-ary search on modern processors. In *Proc. DaMoN*, pages 52–60. ACM, 2009.
6. D. E. Willard. Searching unindexed and nonuniformly generated files in log log N time. *SIAM J. Comput.*, 14(4):1013–1029, 1985.